Scheduling and calibration strategy for continuous radio monitoring of 1700 sources every three days

Walter Max-Moerbeck[a,b]

[a]National Radio Astronomy Observatory, 1003 Lopezville Rd, Socorro, NM, USA
[b]California Institute of Technology, 1200 E California Blvd, Pasadena, CA, USA

arXiv:1409.1859v1 [astro-ph.IM] 5 Sep 2014

ABSTRACT

The Owens Valley Radio Observatory 40 meter telescope is currently monitoring a sample of about 1700 blazars every three days at 15 GHz, with the main scientific goal of determining the relation between the variability of blazars at radio and gamma-rays as observed with the *Fermi Gamma-ray Space Telescope*. The time domain relation between radio and gamma-ray emission, in particular its correlation and time lag, can help us determine the location of the high-energy emission site in blazars, a current open question in blazar research. To achieve this goal, continuous observation of a large sample of blazars in a time scale of less than a week is indispensable. Since we only look at bright targets, the time available for target observations is mostly limited by source observability, calibration requirements and slewing of the telescope. Here I describe the implementation of a practical solution to this scheduling, calibration, and slewing time minimization problem. This solution combines ideas from optimization, in particular the traveling salesman problem, with astronomical and instrumental constraints. A heuristic solution using well stablished optimization techniques and astronomical insights particular to this situation, allow us to observe all the sources in the required three days cadence while obtaining reliable calibration of the radio flux densities. Problems of this nature will only be more common in the future and the ideas presented here can be relevant for other observing programs.

Keywords: active galaxies, , observatory operations, radioastronomy, telescope scheduling

1. INTRODUCTION

Since mid-2007 we have carried out a dedicated long-term monitoring program of blazars at 15 GHz using the Owens Valley Radio Observatory 40 meter telescope.[1] The main goal of this program is to study the location of the gamma-ray emission site with respect to the radio emission site by studying correlated variability. The location of the production site of gamma-ray emission in blazars is not well constrained, and different theories and observations favor a location close to the central engine, while others place it at parsec scales in the radio jet. The goal of our program is to reveal if the radio and gamma-ray emission are correlated, which would be a strong indication of a coespatial origin of the photons for these two energy bands. The monitoring of the gamma-ray band is secured by the continuous operation of the *Fermi Gamma-ray Space Telescope* (*Fermi*) which in normal scanning mode scans the whole sky every three hours.[2] In order to take full advantage of this capability, a matching program in the radio band is highly desirable, which is the purpose of the OVRO 40 meter program.

Here I describe the OVRO 40 meter blazar monitoring program, as well as the observational techniques used for the radio monitoring. The focus of this paper is on the presentation of the method used to schedule the observation of about 1700 sources every three days, to ensure that we obtain proper calibration and reduce the slewing times to maximize the time the telescope spends doing astronomical observations. The ideas presented here could be of use to other monitoring programs confronting similar challenges.

Further author information: Send correspondence to wmax@nrao.edu
Poster presentation at the SPIE Astronomical Telescopes + Instrumentation 2014 in Montréal, Quebec, Canada

2. THE OVRO 40 METER TELESCOPE RADIO MONITORING PROGRAM

The monitoring program started in mid-2007, before the launch of *Fermi*. The original sample of sources consisted of candidate gamma-ray blazars,[3] later augmented by adding the blazars detected by *Fermi* which have high confidence associations with radio blazars.[4,5] We observe all the sources in those catalogs located north of -20° in declination. Total power flux densities are obtained at 15 GHz with a 3 GHz band, and we attempt to observe each source twice per week. A detailed description of the monitoring program can be found in.[1] The program generates high quality radio light curves, as example light curves for our primary flux calibrator and the blazar J2253+1608 show in Figure 1. These light curves are the result of continuous monitoring and the scheduling system has played a significant role in making this program possible.

Our OVRO monitoring program has had a broad impact in the community and has produced a large number of publications by our collaboration and other groups.*

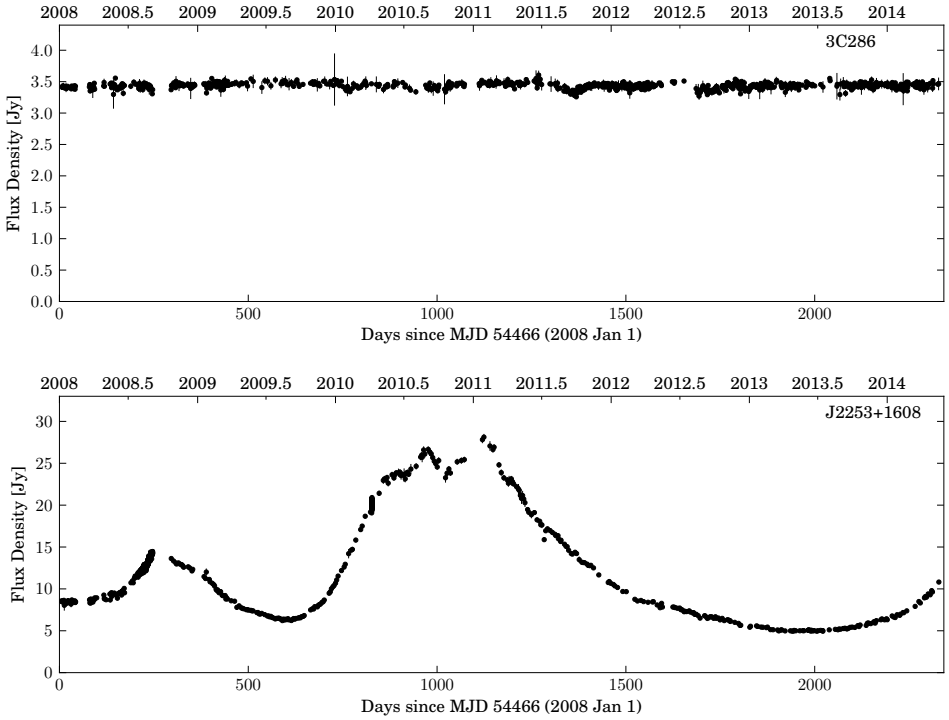

Figure 1. Example radio light curves from the OVRO blazar monitoring program. The upper panel is for 3C 286 our primary flux density calibrator. The lower panel is for the blazar J2253+1608, that shows extreme variability in radio and other bands.

3. SCHEDULING OF THE OBSERVATIONS

The only feasible way to continuously monitor the large number of sources at the cadence required by our program is to make all observations automatic. This demands a scheduling algorithm that can program the observations and calibrations, minimizing the slewing times, and that can run unsupervised. Early in the program we discovered that using the original scheduling and calibration algorithm, some of the sources showed very large flux density variations from day to day, and a subset of them even seemed to have two light curves tracking each other, with one of them at systematically lower flux density levels than the other. The problem was narrowed down to time and sky position variations of the pointing corrections. This section presents a discussion of the requirements and the solution we adopted to solve the scheduling and calibration problem.

*More details about our program and its impact can be found in our website
http://www.astro.caltech.edu/ovroblazars/

3.1 The problem

Due to the large aperture of the telescope and the moderately high brightness of our sources, short integrations provide adequate noise levels in most objects. For the monitoring observations we can obtain a thermal noise level of 4 mJy in just 32 seconds of integration, which becomes 63 seconds when we add the overhead associated with moving the telescope between the 4 segments of the beam switched measurement.[1] A first lower limit in the time required to observe all the sources in the program comes from adding the time it takes to obtain a flux density measurement, which is about $1700 \times 63/3600 \approx 30$ hours. Additionally the time needed to move the telescope from source to source can amount to a significant fraction. A quick estimate can be obtained by assuming all the sources are uniformly distributed in the observable region north of -20° in declination. This comprises an area of 2.7π sr, and a mean distance between sources of 4.5°. The telescope slews at 15° per minute and requires around 12 seconds to settle on each source; taking this into account we can get a lower limit assuming that we can observe all the sources just by jumping into the next one which is at the average distance, in this case we obtain a total slewing time of $1700 \times (12 + 4.5/(15/60))/3600 \approx 14$ hours. To these we have to add the time required to calibrate the flux densities and measure pointing model corrections. Relative flux density calibration requires a calibration diode measurement about once an hour, which adds an approximate overhead of one minute every hour of observations, including slewing. Pointing model corrections are the main restriction as we found that these corrections vary with time and are only valid in regions of less than about 25° in diameter. If we divide the sky into about 100 regions, each one requiring a pointing calibration, which takes about 7 minutes, this amounts to about $7 \times 100/60 \approx 12$ hours. Adding all these numbers we estimate a lower limit on the total observing time of about 60 hours.

An observing sequence in which we visit each source traveling an average distance of 4.5° is far from what really happens, as it ignores the details of the source distribution and its interaction with other observing constraints such as the limits on zenith angle of the observations. The variation of the pointing model corrections limits the possible source arrangements by forcing us to divide the sky into pointing regions of about 25° in diameter, so the observing problem can be separated into two parts: a first level optimization in which the sources are sorted within an observing region, and a second level optimization in which these regions are sorted. The scheduling problem consists in observing all the sources once in each cycle, minimizing the slewing times and respecting the restrictions imposed by the pointing and calibration requirements, and observability of the sources.

This problem is related to a classical optimization problem, the Traveling Salesman Problem (TSP) in which a traveler salesman has to visit a number of cities minimizing the distance traveled. Each city has to be visited only once and the trip ends in the starting city. This problem does not have a known exact solution and requires numerical techniques to search the solution space and find an appropriate solution, which in most cases is only close to the optimum. A direct search of the solution space is only feasible for small number of cities, making the problem hard even for current computational capabilities. This is due to the fact that the number of possible solutions is $N!$, where N is number of cities, a number that becomes very large even for problems of moderate size (e.g., 100 cities). Various methods have been applied to solve the TSP, among them are the nearest neighbor heuristic, simulated annealing, or problem specific heuristics. The problem has found application in many areas of science and engineering including logistics, genome sequencing, electronic circuit manufacturing and many others. In astronomy it has already been used to help in the scheduling of observations of up to a few hundred sources. Large instances of its most simple form have been solved and software that is able to handle them is freely available. A complete and recent review of solution techniques and references for its various application is given in[6†]. Even though the particular example faced by our program is a more complicated version which cannot be handled by these standard tools, some insight can be gained from the existing approaches currently used to tackle this problem.

3.2 A solution

In the particular version of the problem our monitoring program needs to solve there are some additional complications. The first one is that the sources are moving in horizontal coordinates, which are the relevant coordinates for the scheduling problem, and the second one is that the sources can only be visited during a time window in which they are sufficiently high in the sky to reduce the effects of the atmosphere and far from the zenith where the telescope has problems tracking the sources. All these factors make the solution more challenging than the traditional TSP problem. After some experimentation with some of the basic tools used for the TSP, a scheme that solves the practical problem of fitting the sources to the desired 3 day cycle was found and is described below.

[†] A collection of interesting resources related to the TSP can be found in a website maintained by W. J. Cook at `http://www.tsp.gatech.edu/`

The requirement of having pointing model corrections for regions of less than 25° in diameter, requires us to divide the sky into smaller regions. We chose the HEALPix grid to accomplish this as it is widely used in astronomy and has the property of generating regions of equal area and thus a similar number of sources.[7] The main restriction of having a distance between the source and pointing calibrator of less than 25°, can be accomplished by using a grid with 192 pixels over the full sky, each one with $14^\circ.7$ of diameter.[‡] This requirement reduces the complexity of the optimization problem by naturally dividing the problem into two levels, one in which sources are sorted in a region and a second one in which the regions are sorted to make a complete observing cycle.

The first step in the optimization is to assign each source to a region. For each region we need to select a source for pointing calibrations. This source is used to obtain the local pointing model corrections that are later applied to the other sources in the region. A pointing calibrator needs to be bright, unresolved and located in a region free of confusion. Almost all the sources in our monitoring program satisfy the unresolved requirement and are far from the Galactic plane in regions which are usually not affected by confusion; the few which are not suitable are excluded from this procedure. The pointing calibrator is chosen from the sources that are brighter than a certain threshold. Regions with one of the standard flux density calibrators use this for pointing. For other regions, we select among the sources brighter than 400 mJy and choose the one with the minimum average distance to other sources. In regions with no sources brighter than 400 mJy we simply choose the brightest source. For a couple of special cases near the Galactic plane where contamination of the reference fields is severe, sources have been incorporated in nearby regions always ensuring a distance of less than 20° to the pointing calibrator.

For each region we have to solve the problem of visiting the sources starting with the pointing calibrator in an order that minimizes the slewing time. Regions with fewer than 10 sources can be optimized by direct search of the best solution. For regions with at least 10 sources this approach is too costly and we resort to simulated annealing, which can find a good solution in a reasonable time[§]. The extent of this problem can be appreciated by looking at Figure 2, which shows that only a small fraction of the regions can be optimized by direct search.

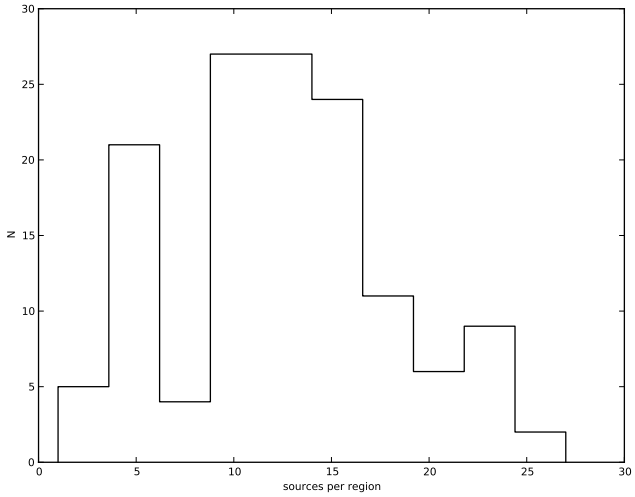

Figure 2. Number of sources per region in the sky. Only regions with fewer than 10 sources can be optimized by direct search; for most of them we use simulated annealing.

Even though we cannot hope to solve the problem exactly with a direct search, we can use simulated annealing, which is an approximation algorithm used to find solutions in optimization problems that do not have exact solutions or when

[‡]This diameter is actually an equivalent angular size given by $\theta_{\rm pix} = \Omega_{\rm pix}^{1/2}$, where $\Omega_{\rm pix}$ is the solid angle subtended by each pixel.

[§]The time it takes for a direct search is proportional to the number of possible solutions. Using our implementation on a desktop computer it takes 108 seconds to test the 9! possibilities for 9 sources. By simply scaling this we can predict that it would take 300 days for 14 sources and 80% of the age of the universe for 22 sources. It was not fun to realize this in my second year of graduate school.

finding such would take an unreasonable amount of time. The basic idea is to simulate the processes of annealing in a metal, in which the material is heated and slowly cooled in order to remove strain and imperfections as a result of minimizing the free energy in the material. In a practical optimization problem, instead of reducing the free energy we are interested in reducing some cost function, and we choose a quantity analogous to the temperature that is a measure of the variability of the cost function. This method is easy to implement and although it does not ensure an optimum solution, it has proven to be efficient at finding good solutions with a reasonable amount of computation.[8,9] In practice we find appropriate solutions in very short times with only a few thousand iterations. The simulated annealing algorithm is summarized below (adapted from[9]). Here a tour is a path that goes from source to source and we denote it by T. The time length is the cost function for the minimization and we called it $C(T)$,

- Compute an initial tour of the sources T and choose an initial temperature $\theta > 0$ and a repetition factor r.
    - As long as the stopping criterion is not satisfied perform the following steps
        - Do the following r times.
            - Perform a random modification of the current tour $T_{\rm mod}$ and compute the time length difference $\Delta = C(T_{\rm mod}) - C(T)$.
            - Draw a uniformly distributed random number x, $0 \leq x \leq 1$.
            - If $\Delta < 0$ or $x < \exp(-\Delta/\theta)$ then set $T = T_{\rm mod}$.
        - Update θ and r
- Output the current tour as solution

This is only a general description of the algorithm that leaves a number of points out. The first one is how to choose a starting temperature θ. There are various recipes for this: one is to explore a number of random paths and use a quantity proportional to the standard deviation on those random tours. In our case we simulate N = 100 random tours and start with $\theta = \theta_i = 3\sigma$, where σ is the variance in those N random tours. Another point is the repetition factor r. After some trial we decided to use r = 100, as no significant improvements were found for larger values. The last point is a rule to update θ and r, which is the so-called cooling schedule. We use the simplest cooling schedule in which r is fixed and θ is reduced by a constant factor, in this case 0.9. The procedure ends when $\theta < 0.01\theta_i$. We initialize the optimization by choosing a random tour which is modified at each step by swapping the order of two sources[¶].

The solution will depend slightly on the elevation for the observations, but the dependence can be ignored and a fixed value obtained at 45° of elevation can be used for the following stage which is described below. The sample path on the sky for one of the regions is illustrated in Figure 3 for a region with 21 sources.

The second stage consists of sorting the regions. This has the extra complication that the regions are only observable for a limited time and that calibrators have to be observed every day. Several approaches were tested, including sorting by declination and right ascension, or by nearest neighbor, but none of them was able to accommodate all the sources in the required 3 day cycle. After some experimentation it was found that a heuristic approach that starts by giving higher priority to observations of southern sources and then moves slightly to the north is able to fit all the sources in the 3 day cycle. This heuristic approach is motivated by the fact that southern sources have a very limited observing window while circumpolar sources are observable at any time. The heuristic approach effectively uses circumpolar sources to fill gaps in the schedule where no other regions are observable. An example path through all the 136 regions with sources is shown in Figure 4.

This system has been in use since March 2009. It has allowed us to reduce large systematic errors associated with inadequate pointing model corrections for large regions of the sky, a problem that affected a fraction of our early light curves. All those problematic data points have been eliminated and we improved the reliability of our measurements and increased the effective cadence on each source. Although this original heuristic approach has been revised since its adoption we still use the software infrastructure developed in 2009, which has proven to be flexible enough to accommodate new developments.

[¶]In the context of the TSP this is the *city-swap heuristic*

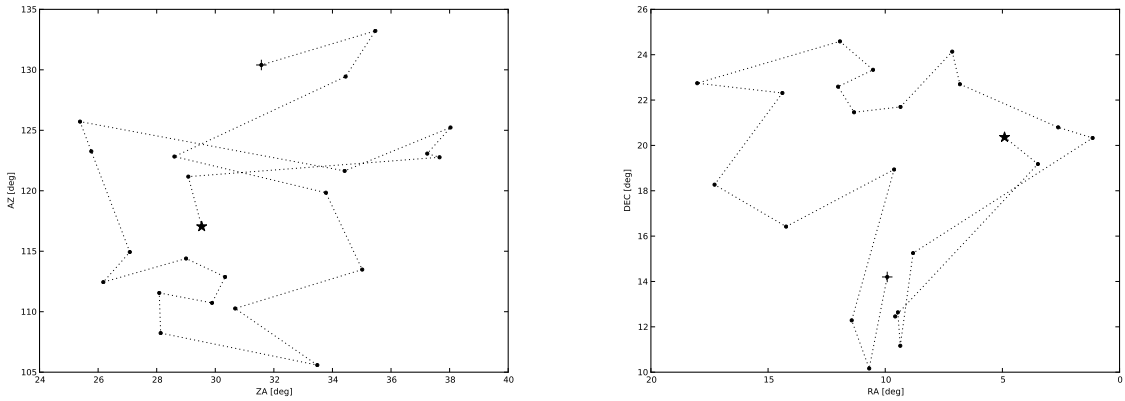

Figure 3. Sample path of the telescope for one of the regions in horizontal coordinates (left panel) and equatorial coordinates (right panel). The first source is marked with a star and the last one with a cross. There are a total of 21 sources in this region.

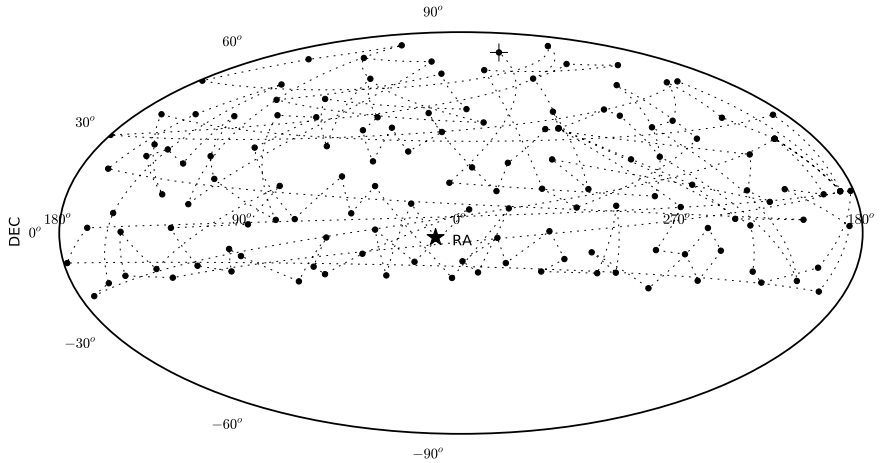

Figure 4. Sample telescope path in equatorial coordinates for a full observing cycle of three days. The first region is marked with a star and the last one with a cross.

4. SUMMARY

This paper presents the implementation of a practical solution to the telescope scheduling, calibration, and slewing time minimization problem as applied to the OVRO 40 meter telescope blazar monitoring program. This solution combines ideas from optimization with astronomical and instrumental constraints. Problems of this nature are expected to be more common in the future and the ideas presented here can be relevant for other observing programs.

The Owens Valley Radio Observatory 40 meter telescope has been continuously observing blazars since mid-2007. In that time, and enabled in part by the developments presented in this paper, it has been producing high quality light curves used in many refereed publications. Thanks to these data our group has provided insight into the relation between the radio and gamma-ray emission in blazars,[10,11] the amplitude of the radio variability and its relation to source class,[1,12,13] and the correlations between the radio and gamma-ray variability as a tool to locate the gamma-ray emission site.[14] More than 40 peer reviewed publications, several conference proceeding and Astronomer Telegrams have used our data for multi-wavelength studies of blazars and other variable sources.

ACKNOWLEDGMENTS

This work was done while I was a graduate student at Caltech. The successes of the OVRO 40 meter blazar monitoring program are the result of the effort of many people. I thank all the members of the blazar monitoring team at Caltech and elsewhere, for insightful conversations about radio astronomy observations, data analysis techniques and the physics of active galactic nuclei. Particular acknowledgments go to Joseph Richards, Talvikki Hovatta, Anthony Readhead and Tim Pearson, who are the core of the group. I thank Russ Keeney for his support at OVRO. The OVRO program is supported in part by NASA grants NNX08AW31G and NNX11A043G and NSF grants AST-0808050 and AST-1109911. Support from MPIfR for upgrading the OVRO 40-m telescope receiver is acknowledged. The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

REFERENCES

[1] J. L. Richards, W. Max-Moerbeck, V. Pavlidou, O. G. King, T. J. Pearson, A. C. S. Readhead, R. Reeves, M. C. Shepherd, M. A. Stevenson, L. C. Weintraub, L. Fuhrmann, E. Angelakis, J. A. Zensus, S. E. Healey, R. W. Romani, M. S. Shaw, K. Grainge, M. Birkinshaw, K. Lancaster, D. M. Worrall, G. B. Taylor, G. Cotter, and R. Bustos, "Blazars in the Fermi Era: The OVRO 40 m Telescope Monitoring Program," *ApJS* **194**, p. 29, June 2011.

[2] W. B. Atwood, A. A. Abdo, M. Ackermann, W. Althouse, B. Anderson, M. Axelsson, L. Baldini, J. Ballet, D. L. Band, G. Barbiellini, and et al., "The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission," *ApJ* **697**, pp. 1071–1102, June 2009.

[3] S. E. Healey, R. W. Romani, G. Cotter, P. F. Michelson, E. F. Schlafly, A. C. S. Readhead, P. Giommi, S. Chaty, I. A. Grenier, and L. C. Weintraub, "CGRaBS: An All-Sky Survey of Gamma-Ray Blazar Candidates," *ApJS* **175**, pp. 97–104, Mar. 2008.

[4] A. A. Abdo et al., "The First Catalog of Active Galactic Nuclei Detected by the Fermi Large Area Telescope," *ApJ* **715**, pp. 429–457, May 2010.

[5] M. Ackermann et al., "The Second Catalog of Active Galactic Nuclei Detected by the Fermi Large Area Telescope," *ApJ* **743**, p. 171, Dec. 2011.

[6] D. Applegate, R. Bixby, V. Chvátal, and W. Cook, *The Traveling Salesman Problem: A Computational Study*, Princeton Series in Applied Mathematics, Princeton University Press, 2011.

[7] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, "HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere," *ApJ* **622**, pp. 759–771, Apr. 2005.

[8] P. van Laarhoven and E. Aarts, *Simulated Annealing: Theory and Applications*, Ellis Horwood series in mathematics and its applications, Springer, 1987.

[9] G. Reinelt, *The traveling salesman: computational solutions for TSP applications*, Lecture notes in computer science, Springer-Verlag, 1994.

[10] M. e. a. Ackermann, "The Radio/Gamma-Ray Connection in Active Galactic Nuclei in the Era of the Fermi Large Area Telescope," *ApJ* **741**, p. 30, Nov. 2011.

[11] V. Pavlidou, J. L. Richards, W. Max-Moerbeck, O. G. King, T. J. Pearson, A. C. S. Readhead, R. Reeves, M. A. Stevenson, E. Angelakis, L. Fuhrmann, J. A. Zensus, M. Giroletti, A. Reimer, S. E. Healey, R. W. Romani, and M. S. Shaw, "Assessing the Significance of Apparent Correlations between Radio and Gamma-Ray Blazar Fluxes," *ApJ* **751**, p. 149, June 2012.

[12] T. Hovatta, V. Pavlidou, O. G. King, A. Mahabal, B. Sesar, R. Dancikova, S. G. Djorgovski, A. Drake, R. Laher, D. Levitan, W. Max-Moerbeck, E. O. Ofek, T. J. Pearson, T. A. Prince, A. C. S. Readhead, J. L. Richards, and J. Surace, "Connection between optical and γ-ray variability in blazars," *MNRAS* **439**, pp. 690–702, Mar. 2014.

[13] J. L. Richards, T. Hovatta, W. Max-Moerbeck, V. Pavlidou, T. J. Pearson, and A. C. S. Readhead, "Connecting radio variability to the characteristics of gamma-ray blazars," *MNRAS* **438**, pp. 3058–3069, Mar. 2014.

[14] W. Max-Moerbeck, T. Hovatta, J. L. Richards, et al. *submitted* , 2014.